




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 13/14		A1	(11) International Publication Number: WO 00/17765
			(43) International Publication Date: 30 March 2000 (30.03.00)
(21) International Application Number: PCT/US99/21248		(81) Designated States: AU, BR, CA, CN, CZ, HU, ID, IL, IN, IS, JP, KR, LT, LV, MX, NZ, PL, RO, RU, SG, TR, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 22 September 1999 (22.09.99)			
(30) Priority Data: 9803246-9 24 September 1998 (24.09.98) SE		Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	
(71) Applicant (for all designated States except US): MIRROR IMAGE INTERNET, INC. [US/US]; Suite 4800, 18 Commerce Way, Woburn, MA 01801 (US).			
(72) Inventor; and			
(75) Inventor/Applicant (for US only): LINDBO, Sverker [-SE]; Bjorkliden 16, S-187 41 Täby (SE).			
(74) Agent: PAYNE, R., Thomas; Cummings & Lockwood, Four Stamford Plaza, Stamford, CT 06904 (US).			

(54) Title: AN INTERNET CACHING SYSTEM AND A METHOD AND AN ARRANGEMENT IN SUCH A SYSTEM



(57) Abstract

The present invention refers to an Internet caching system and to an arrangement and a method for serving request for Internet information files in an Internet caching system. The system is built as a two tier caching system. In order to decrease the load on a central cache server (130), an intermediate arrangement (110) interconnects the local servers (100) of the system to the central cache server (130). This arrangement communicates with the local cache servers in accordance with a protocol used for communicating between cache servers. When requesting an Internet information file from the central cache server, the arrangement uses the Structured Query Language. Thus, the central cache server (130) is primarily devoted to answer plain SQL queries.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Larvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

AN INTERNET CACHING SYSTEM AND A METHOD AND AN
ARRANGEMENT IN SUCH A SYSTEM

Technical field of the invention

5 The present invention refers to an Internet caching system and to an arrangement and a method for serving requests for Internet information files in an Internet caching system.

Background art

10 The Internet and its currently most used feature, the World Wide Web (WWW), has in recent years developed into an enormous source of information. Anybody can provide any information, such as text, pictures, audio and video, on the World Wide Web where it can be easily retrieved by users anywhere in the world as long as they 15 have access to the Internet.

20 The major problem facing the Internet is the growing demand for communication capacity as users access information from anywhere in the world. It is estimated that the World Wide Web traffic already exceeds all conventional telephone and facsimile traffic on most international communication lines. More transmission and switching capacity is continuously added, but it is a slow and expensive process and demand continues to outstrip supply.

25 The content of the World Wide Web is getting to be unmeasurable and probably comprises several hundreds of Terabytes (as of summer 1998). However, a relatively small subset of all this information accounts for a huge portion of the information actually being accessed.

30 Therefore, in order to minimize bandwidth used and latency involved when accessing information on the Internet, different caching techniques are currently in use for limiting the amount of information that has to be

transferred over the Internet and for limiting the distance over which the information is transferred.

In the field of caching WWW objects, or Internet information files, there are basically two approaches, 5 client side caching and server side caching. The simplest form of client side caching is virtually used by every WWW browser today. The browser retains a cache on the user's computer with the last accessed Internet information files. When the user for a second time wishes to 10 access a particular information file, the browser retrieves it from its cache rather than making a request for it over the Internet.

In order to help a neighboring user, a proxy server caching method, another form of client side caching, can 15 be used. In this scheme a cache is placed at a WWW proxy node to which a number of neighboring users are connected, such a proxy node could for example be a server located at a company. When a WWW client wants to access a WWW server on the Internet, the client sends a http 20 request to the proxy node, or WWW proxy server, rather than sending it directly to a server on the global Internet. Instead it is the proxy server that sends the request to a WWW server on the global Internet, caches the response and returns the response to the client. 25 Thus, the first time an information file is requested it is transferred over the Internet and stored in the cache of the WWW proxy server. Subsequent requests for the same information file from any client connected to the WWW proxy server can then be resolved locally, rather than 30 making http requests to a WWW server over the global Internet. Proxy server caching can also be used outside the premises of a company, or some other organization, by implementing the scheme described above at a regional Internet cache server to which a number of clients are 35 directly or indirectly connected.

Depending on the size and homogeneity of a user community using a cache at a server, about 20-40 Giga-

bytes of cache storage will (spring 1998) reduce the Internet traffic generated by the user community by 30-50 %. As the growth of the information provided by the Internet and the WWW continues, it is highly likely that 5 the required cache size will have to increase over time to retain the hit rate, i.e. the proportion of the information files requested that are transferred from the cache server. Furthermore, it would give significant benefit for the performance and utilization of the 10 Internet if the hit rate could be increased to 75% or more. With the typical end user behavior, this would require a much larger cache, currently in the order of 200-400 Gigabytes, but also require very many members in the end user community, currently several hundreds of 15 thousands. The reason is that the larger the end user community, the larger the probability that someone else within the community has previously accessed a requested file, especially if the users share some common interest.

Installing a large cache is easily achieved by 20 acquiring the appropriate computer and the appropriate disk capacity. However, it is also required that the cache is able to handle all requests from the participating end users. Using current technology, it is not possible for one single processor computer to serve the 25 requests from several hundred thousand end users. Hence, several systems have been presented to deal with this problem, here outlined under the names of their major proponents.

Cisco Systems, Inc. proposes that the end users are 30 connected to a backbone router which is programmed to transparently redirect all WWW requests to a group, or "Farm" of dedicated cache appliances, or "Cache Engines". Each Cache Engine handles a subset of all origin WWW servers, based on grouping of the IP (Internet Protocol) 35 addresses. The solution scales up to 32 Cache Engines in parallel, which corresponds to serving approximately 500.000 subscribing end users.

Inktomi Corporation suggests that a switch, a so called layer 4 switch, is used to redirect all requests for WWW pages to an "Inktomi Traffic server". A cluster of powerful computers are used, which all share the same 5 disk storage system. This solution scales up to 16 parallel workstations, which also corresponds to about 500.000 subscribing end users. However, having several computers accessing the same disk storage system adds complexity and requires management, i.e. some of the 10 capacity of each computer is not available for processing requests.

Network Appliance, Inc. proposes a two tier caching solution. The system has several local caches near the end users. These local caches communicate with a central 15 cache using the Internet Cache Protocol (ICP) when a cache miss occurs at the local level. If the requested file is present in the central cache, it will be transferred to the local cache and then forwarded to the end user. If the requested file is not in the central cache 20 either, the central cache will make a request to the origin server and forward the file to the local cache, which in turn forwards the file to the end user. The central cache thus handles ICP requests from the local caches and communicates with the origin server in the 25 case of a cache miss at the central cache. For scalability, there can be several central caches in parallel, each handling a subset of the origin servers. This means that the local caches are able to address each request to the correct central cache server. Since this protocol is 30 not standardized, it means that all local caches have to be delivered from Network Appliance, Inc.

All of these solutions have the drawback that a central cache server needs to handle extensive communication in one way or another. This results in low utilization of the server's capacity and difficulties in 35 serving hundreds of thousands users, which is required in order to obtain a high hit rate. By adding more servers,

the systems are made more expensive and more complex. The complexity of the system adds to the overhead and, hence, to a low utilization of the relatively expensive resources that the servers represent.

5

Summary of the invention

An object of the present invention is to overcome the drawbacks with the presently known techniques for caching information files on the Internet and to provide 10 a solution for caching information files in a cost-effective way.

Another object of the present invention is to provide a solution for how user's requests for cached information files are to be served by a caching system in a 15 fast and cost-effective way.

Yet another object is to provide a cache server solution which is able to cope with the growing numbers of information files being provided by the Internet and the World Wide Web.

20 Yet another object is to provide a solution for obtaining a high hit rate percentage for information file requests directed to a caching system with a minimum of cost.

25 Yet another object of the present invention is to provide a scaleable caching system which is scaleable in a standardized way.

The above objects are achieved by an Internet caching system and a method for serving requests for Internet information files in an Internet caching system 30 in accordance with the appended claims.

According to a first aspect of the invention, there is provided a method for serving requests for Internet information files in an Internet caching system, which method comprising the steps of receiving, at a local 35 Internet cache server, a user request from a user for an Internet information file; in response to the received request, making a query for said information file, if

said information file has not been cached by said local server; in response to a reply to said query, making a file request for said information file, wherein said file request is directed to a feeder means if said reply 5 indicates that a central file server, storing cached Internet information files, has said information file cached; and querying, from said feeder means in response to said file request, said central file server for said information file, in order to decrease the load on said 10 central file server.

According to a second aspect, there is provided an arrangement in an Internet caching system, said system comprising at least one local cache server and at least one central file server, both of which servers stores 15 cached Internet information files, which arrangement, for decreasing the load on said central file server, includes a Feeder communicating with said local cache server and with said central file server, wherein said Feeder includes first means for receiving a request for an 20 Internet information file from said local cache server; second means for deriving a query from an alphanumerical string received from said local cache server; and third means for querying said central file server for said Internet information file using said query derived by 25 said second means.

According to a third aspect, there is provided an Internet caching system, which system comprises a set of local Internet cache servers, wherein each local cache server is arranged to receive requests from users for 30 Internet information files; at least one central file server included in a central cache site and storing cached Internet information files; and feeder means interconnecting said set of local cache servers with said central file server, said feeder means including at least 35 one Feeder, which Feeder comprises means for communicating with at least one local cache server in accordance with a protocol used for communicating between Internet

cache servers and means for retrieving Internet information files from said central file server using data base queries, thereby decreasing the load on said central file server.

5 The invention is based upon the idea of connecting a number of dedicated computers to a central file server, or central cache server, storing Internet information files. Relative to the central cache server, these additional computers are low end computers. The dedicated
10 computers are arranged to decrease the load on the central cache server by performing some of the tasks normally handled by the central cache server itself. In this way the central cache server is able to serve the local cache servers connected to the central server, or
15 rather connected to the central server via the dedicated computers, in a fast and cost-effective way. Maximum use is made of the expensive hardware forming the actual central file server and its file repository in which the files are cached, while specialized inexpensive machines
20 around the file server perform time consuming and time critical tasks in parallel.

Thus, the inventive feeder means, or Feeders, are realized by machines being separate from any machine realizing a central file server. This will decrease the load on the central file server, which then is able to dedicate more processing time to the actual retrieval of cached information files. Hence, the central file server is able to serve a large community of users in an efficient way. Since user requests, via requesting local
30 cache servers, are served more effectively, the number of user requests served can be increased, which in turn enables the central file server to obtain a higher hit rate percentage for its cache.

According to an embodiment of the present invention,
35 the feeder means communicates with the local cache servers, on behalf of the central file server, in accordance with a protocol used for communicating between

Internet cache servers. The currently used protocol is either the Internet Cache Protocol (ICP) or the Cache Digest, but could be any other conventional or future protocol used for the same purpose. Thus, by placing the 5 task of accepting, and replying to, queries and/or requests for information files in machines being separate from the central file server machine, the load on the central file server is decreased considerably.

When a local cache server receives a request from a 10 user for an information file, which file has not been cached at the local server, the local server starts with making a query for that file. In one embodiment the query is directed to a table, or data base, being internal to, or directly connected to, the local server. If said table 15 indicates that the queried file is cached by the central file server, the local server will request the file from the feeder means, or Feeder. This querying and requesting is then preferably performed in accordance with the Cache Digest protocol. However, as with the request from the 20 user to the local server, the request from the local server to the Feeder may be in accordance with any layer three protocol, for example an HTTP request.

In another embodiment, the query from the local server is directed to the Feeder. Included in the query, 25 for example an ICP query, is the URL of the queried information file. The Feeder derives a query number from the alphanumerical URL of the received query for an information file, which query number then is used by the Feeder for querying the central file server for the 30 information file. The Feeder queries the file server for information files using a standard SQL query (Structured Query Language). If the queried file is present at the central file server, i.e. if there is a cache hit, the queried file is transferred from the central server, via 35 the Feeder, to the local server. To have the central file server initiate a file transfer as an answer to a SQL query, rather than as an answer to a query, such as an

ICP query, from the local cache server, means considerable capacity savings at the central file server.

Alternatively, the query number is derived from said alphanumerical URL and from a part of a header information included in said query. This part of the header information contains specific user information of the original requester, for example, the language he is using, enabling the central file server to respond in accordance with this specific information. The query number corresponding to an information file is derived by 10 using any hash algorithm, preferably using an MD5 hash algorithm.

In the embodiment where the local server makes an internal query for the information file, the Feeder derives the query number from the following request 15 directed to the Feeder by the local server. The alphanumerical string used for deriving the query number is the string included in said request, for example the URL of an HTTP request. The query number is then used by the Feeder when querying the central file server for the 20 information file, preferably using an SQL query. Again, it is advantageous to also include at least part of an header information field of said request as the basis for deriving said query number.

In order to decrease the load on the central file 25 server even further, the Feeder preferably includes a table storing information relating to each information file being cached by the central file server. The table, for example, being a memory resident MD5 indexed hash 30 table. By searching said table, the Feeder can conclude whether or not a queried information file is cached by the central file server, without having to query the server, and, hence, a faster reply may be given by the Feeder to the query from a local server.

According to another embodiment of the present 35 invention, the Internet caching system further comprises updater means, or an Updater, for updating the set of

information files being cached by the central file server. The updating procedure consists of transferring a copy of a file cached at a local server to the central server. The transferred file is a file which, as a consequence of a cache miss at the central server when querying for the file, has been retrieved from its origin server by the local server and then been cached by the same.

Thus, the central file server, or central cache server, does not itself retrieve a non-cached file and is therefore not burdened with having to make a request to an origin server for a file because of a cache miss when serving a local cache server. Instead, when the Feeder evaluates a query from the local cache server for an information file, and concludes that the queried file is not cached at the central file server, the Feeder directs a reply to the querying local server, indicating that the file is not available, and then orders the Updater to update the central file server. Upon reception of the reply, which thus indicates a cache miss, the local cache server retrieves the file in question from its origin server. Upon reception of the order to update the central file server, the Updater requests a copy of the file from the local server and transfers the thereby received file copy to the central cache server where it is stored. The transferring and storing procedure is preferably performed at a time when the overall load on the central file server is low, and when the local server has been given enough time to retrieve the file from its origin server.

However, should the local server be located behind a firewall, the Updater will request a copy of the file from its origin server, which copy then is stored in the central cache server. In this case, it is preferred that the Feeder does not order the Updater to commence the updating procedure until after a certain number of queries for the same particular information file have

been received, where these queries originate from local servers being located behind firewalls. Preferably, the Updater is realized by a machine being separate from the machines realizing the Feeders, as well as being separate from any file server machine. This is an advantage since file requests, for example HTTP requests, to origin servers may take unpredictable amounts of time and thus lead to an unpredictable load on the machine performing the requests. However, in a simplified system it is possible to realize the Updater in the same machines as those which realize the Feeders, while still being separate from any central file server machine. In an embodiment where the machines implementing the Updater and the Feeders interconnects the local cache server with the central file server, without the machines themselves being included in the central cache site together with the central file server, the separation of these machines from the central file server machine is evident.

Certain Internet information files are not suitable for caching. Such files are sometimes called dynamic information files, the term dynamic comes from that these files are continuously updated at the origin server, examples of such files are files with stock quotes, weather reports and so on. One preferred way of treating the existence of dynamic files is to uphold a list of known uncachable files in either the Updater or in the local servers. In this way the communication in the system, as a result of a user requesting such a file, can be minimized.

According to yet another embodiment of the present invention, several central file servers are included in a central cache site, each file server caching information files associated with original host names, IP-addresses or derived query numbers, within a defined range. Based upon either the original host name, IP-address or the derived query number of a requested information file, the Feeder addresses the query to the file server caching

files in the appropriate range. In this scaleable solution each file server has its own disk system, thus minimizing overhead. Furthermore, the central cache site is scaleable with third party file servers because of the 5 standardized protocols used by the site.

In order for the communication between the central file server and the low end computers, i.e. the Feeders and the Updaters, to be fast, each low end computer is preferably connected to the central file server by means 10 of a dedicated wire, alternatively, if there are several file servers, by means of a dedicated network. This network is either a private or a public network. In the latter case, at least part of the network capacity is preferably reserved for the communication in question. 15 The network used can, of course, also be a part of the Internet, also in a non-dedicated way. The type of connection used between the central file server and the low end computers is very much dependent upon where the low end computers, or Feeders and Updaters, are located, at 20 the same site as the central file server, or, at a location being different from the location of the central file server.

Moreover, it is preferred that the central cache site serves a defined set of local cache servers, which 25 set in turn serves a linguistically and culturally homogenous user community. This will further increase the hit rate percentage at the central cache level since it is more likely that the same information files are requested more than once.

30 Using the present invention, an operator of an Internet caching system, handling information file requests in accordance with the present invention, is able to provide a fast, cheap and effective way of serving a large number of subscribing customers. These 35 customers preferably being different Internet service providers, companies or other organizations connected to the inventive central cache site, or inventive Feeders/

Updaters, with their own local cache servers, or, being connected as clients to a system encompassing the whole inventive caching system formed by the central cache site, including Feeders and an Updater, and its connected local cache servers. Of course, a customer may very well also be a single user constituting a single WWW client connected directly to the inventive system. Also, a large company or Internet service provider can choose to operate the inventive system on its own rather than being connected to such a system being operated by another party. Furthermore, since the inventive caching system is built around standardized protocols, such as ICP and SQL, local cache servers and central file servers from any manufacturer can be included in the system as long as these protocols are supported.

Within the scope of the present invention a local Internet cache server is to be interpreted as a proxy node, preferably a WWW proxy node, retaining a cache for the users, or WWW clients, connected to the proxy node.

Items cached on a local Internet cache server or a file server at a central cache site are any non-dynamic files which are accessible using the Internet and containing any type of information. Thus, a number of different type of files and different namings of such files are included by the term Internet information file used in the present invention, such as binary, text, picture, audio and video files, HTTP (HyperText Transfer Protocol) files, WWW files, FTP (File Transfer Protocol) files, WWW pages, WWW objects, and so on. Besides files being accessed using the HTTP or FTP protocol, any file being accessed over the Internet in accordance with any layer 3 protocol is also included by the term Internet information file. A further example of a protocol that can be used is the WTP protocol (Wireless Transport Protocol) used within the WAP (Wireless Application Protocol) standard.

According to a fourth aspect of the present invention, the invention encompasses a computer-readable medium, on which is stored one or several computer programs of instructions for one or several general purpose computers, comprising means for enabling said one or said several computers to perform the steps disclosed in the appended claims 1 - 17.

According to a fifth aspect of the present invention, the invention encompasses one or several program storage devices containing one or several sequences of instructions, for one or several general purpose computers, for performing the steps disclosed in the appended claims 1 - 17.

The above mentioned and further aspects and features of, as well as advantages with, the present invention, will be more fully understood from the following description, with reference to the accompanying drawings, of exemplifying embodiments thereof.

20 Brief description of the drawings

Exemplifying embodiments of the invention will now be described below with reference to the accompanying drawings, in which:

Fig. 1 schematically shows an embodiment of an Internet caching system according to the present invention;

Fig. 2 schematically shows another embodiment of an Internet caching system according to the present invention;

Fig. 3 schematically shows a flow chart of the operations performed by a local cache server in Fig. 2;

Fig. 4 schematically shows a flow chart of the operations performed by a Feeder in Fig. 2;

Fig. 5 schematically shows a flow chart of the operations performed by an Updater in Fig. 2; and

Fig. 6 schematically shows yet another embodiment of an Internet caching system according to the present invention.

5 Detailed description of preferred embodiments

With reference to the block diagram shown in Fig. 1, an embodiment of the present invention will be described. In Fig. 1 a number of local cache servers 100 are shown. These local servers 100 are, via the Internet, connected 10 to feeder means 110, here exemplified with a Feeder 110. The number of Feeders 110 and the number of local cache servers 100 indicated in Fig. 1 is merely an example, and the embodiment is not restricted to these numbers.

However, regardless of the number of Feeders, each 15 Feeder is in this embodiment connected to one single central file server. In Fig. 1 Feeder 110 is connected to a central file server 130. This central file server includes a storage medium (not shown) on which Internet information files are stored, i.e. cached, and is implemented by a high end computer, such as a Sun Ultra Sparc or DEC Alpha Computer. Each Feeder 110 on the other hand 20 is implemented by a low end computer, such as a conventional Personal Computer, and constitutes a front end machine which handles the communication between the local 25 cache servers 100 and the central file server 130.

The Feeder 110 communicates with the local cache servers 100 using the Internet Cache protocol, which is a message based protocol used for communicating between cache servers over the Internet. Hence, the Feeder 110 30 replies to an ICP query for a cached Internet information file, the query being received from one of the local cache servers 100, with an ICP reply. This ICP reply indicates either a cache hit (ICP_OP_HIT) or a cache miss (ICP_OP_MISS).

35 In accordance with the Internet Cache Protocol, the ICP query received by the Feeder includes the URL of the queried information file. From this URL the Feeder 110

derives a query number, corresponding to the queried information file, using an MD5 hash algorithm. Using the query number, a memory resident MD5 indexed hash table 115 is then searched. Included in the Feeder 110 is a RAM (Random Access Memory) 116 in which the indexed table is stored. The indexed table 115 comprises an entry for each query number corresponding to an Internet information file cached at the central file server 130. Searching the indexed table 115 comprises searching the entries for a query number matching the derived query number. If a matching query number is found in the table, this is an indication that the queried information file is cached by the central file server 130, and, as a consequence, the ICP reply to the local server 100 will indicate a cache hit. Correspondingly, if a matching query number is not found in the table 115, this indicates that the queried information file is not cached by the central file server 130, and, as a consequence, the ICP reply will indicate a cache miss.

The means for deriving the query number using the MD5 hash algorithm and for searching the indexed table is a microprocessor 120, together with an appropriate software module, included in the Feeder 110. The microprocessor executes the software module, which execution results in derived query number and in a search of the indexed table 115. The implementation of this software module is straight forward for a person skilled in the art of programming.

If the reply from the Feeder 110 to the local server 100 indicates a cache hit, the local server will request the information file from the Feeder using the HyperText Transfer Protocol, which is a protocol used for accessing WWW objects over the Internet. That is, an HTTP request is transmitted to the Feeder, which request includes the URL of the requested file.

When communicating with the central file server 130, the Feeder 110 uses common SQL queries. Upon reception of

the HTTP request, the Feeder will retrieve the query number which was previously derived from the URL of the corresponding ICP query. Alternatively, the URL of the HTTP request is used for once again deriving the query 5 number. The Feeder then uses the query number in a standard SQL query directed to the central file server. As a response, the central file server 130 will transfer the information file in question to the Feeder 110, which in turn transfers the information file to the local 10 server 100 that issued the request for the file.

If the reply from the Feeder 110 to the local server 100 indicates a cache miss, the local server will make an HTTP request to the origin server (not shown) of the requested file, cache the then received file and transfer 15 a copy of the file to the requesting user (not shown).

The means for implementing the execution of the Internet Cache Protocol in the Feeder 110 is the microprocessor 120 included in the Feeder. The microprocessor also implements the means for receiving an HTTP request 20 from the local server 100 as well as the means for querying the central file server 130 using the SQL. The operations to be performed by the microprocessor are controlled by appropriate software modules, being part of the abovementioned means. The implementation of these 25 software modules would be well known for a person skilled in the art of programming and being familiar with the protocols in question.

Another embodiment of an Internet caching system according to the invention is described with reference to 30 Fig. 2. The system in Fig. 2 differs from the one shown in Fig. 1 in that the Internet caching system comprises an Updater 240, i.e. updater means, being connected to the central file server 230, the Feeder 210 and, via the Internet, the local cache servers 200. Thus, Fig. 2 35 illustrates the inventive arrangement encompassing an Updater 240 as well as a Feeder 210.

Besides what is being described below regarding the elements in Fig. 2, the elements of Fig. 2, having corresponding elements in Fig. 1, operate and interact in accordance with what has previously been described with reference to Fig. 1. Therefore, only the features of these elements being of relevance to the embodiment illustrated by Fig. 2 are described below.

The Updater 240 is responsible for updating the storage medium (not shown) associated with the central file server 230 with new cached information files. As described with reference to Fig. 1, when the local server 200 receives a cache miss in an ICP reply from the Feeder 210, as a response to a previous ICP query to the same, the local server 200 makes an HTTP request for the file to its origin server (not shown). The requested file is then received and cached by the local server 200. After a predetermined time, as a consequence to the reported cache miss in the ICP reply, the Feeder 210 will instruct the Updater 240 to update the central file server.

The Updater 240 receives, from the Feeder 210, the URL of the queried file and the identity of the local server 200 which queried for the file. An HTTP request for the file is then made from the Updater to the specific local server. Upon reception of the requested file, the Updater stores, i.e. caches, the file at the central file server 230. When the file has been stored, the Updater instructs the Feeder to add the query number corresponding to the file in question in the indexed table 215 stored in the RAM area 216.

The means for requesting the information file from the local cache server 200 and the means for caching the received information file at the central file server 230 is a microprocessor 260, together with appropriate software modules, included in the Updater 240. The implementation of these software modules would be well known for a person skilled in the art of programming.

An example of the operations performed by a local cache server 200 in the embodiment of Fig. 2 will now be described with reference to the flow chart in Fig 3.

In step 300 the local cache server 200 receives a 5 request for an Internet information file from a client- served by the particular local cache server. However, the file request may also be received from the Updater 240, which operates in accordance with the description referring to Fig. 5. The local cache server then in step 10 301 searches its locally cached files for the file requested. If it finds the file, the file is transferred to the requesting client or to the Updater 240, this is indicated with step 302.

If the local cache server 200 does not find the 15 requested file during the search, i.e. it has not cached the requested file, it examines in step 303 if the request came from the Updater. If this condition is true, a message is returned to the Updater in step 304 indicating that the requested file is not available. If the 20 conditional step 303 is false, i.e. if the request came from a client, an ICP query is sent in step 305 to the Feeder 210. In the next step 306, the local cache receives an ICP reply from the Feeder 210 indicating whether or not the central file server 230 has the 25 requested file cached. In step 307 the ICP reply is evaluated. If the reply indicates a cache miss, i.e. the requested file was not centrally cached, the local cache server 200 makes a HTTP request for the file directed to the origin server of the file. If the reply on the other 30 hand indicates a cache hit, the local cache makes a HTTP request to the Feeder 210 for the file, this is indicated with step 309. In the next step 310, the local cache server receives the requested file from the Feeder. Finally, in step 311, the file is transferred to the 35 client which requested the file.

The operations performed by the Feeder 200 in the embodiment of Fig. 2 is now described with reference to the flow chart in Fig. 4.

In step 400 the Feeder 210 receives an ICP query 5 regarding an Internet information file from any of the local cache servers 200 being handled by the Feeder. The query includes the URL of the queried information file. From this URL the Feeder 210 in step 401 derives a query 10 number using an MD5 hash algorithm, which query number is used in step 402 when searching an indexed MD5 hash table being resident in the memory 216 of the Feeder 210.

If the number is not found during the search in the hash table, the Feeder in step 403 sends an ICP reply indicating a cache miss back to the local cache server 15 200 from which the ICP query was received. In step 404 the Feeder 210 then orders the Updater 240 to retrieve the non-cached queried file by passing the URL of the queried file to the Updater. In step 405 the Feeder 210 adds the query number corresponding to the queried file 20 in the indexed hash table 215. This is done in response to that the Updater 240 indicates to the Feeder that the queried file has been transferred from the local server 200 and stored in the central file server 230. The operation of the Updater 240 will be further described 25 with reference to Fig. 5.

If the Feeder 210 in the conditional step 402 finds the query number during the search in the hash table 215, it will in step 406 send an ICP reply indicating a cache hit back to the local cache server 200 from which the ICP 30 query was received. In step 407 the Feeder then receives an HTTP request from the local cache server 200 which previously issued the ICP query. Similar to the ICP query, the HTTP request includes the URL of the requested information file. In step 408 the Feeder 210 retrieves 35 the previously derived query number corresponding to the file. With this query number the Feeder in step 409 queries the central file server 230 for the requested

information file using a standard SQL query. In step 410 the Feeder as a response receives the cached information file from the central file server 230, and in the next step 411, the requested cached Internet information file 5 is transferred from the Feeder 210 to the requesting local cache server 200.

The operations performed by the Updater 240 in the embodiment of Fig. 2 is now described with reference to the flow chart in Fig. 5.

10 In step 500, the Updater 240 receives an order from the Feeder 210 indicating that a particular file should be requested. The file in question was previously requested by the local cache server 200, but the Feeder found that the central cache server 230 had not cached 15 the file. The order includes the URL of the file as well as the address of the local cache server 200 which requested the file from the central cache 230. The Updater will then, in step 501, check the requested file of the order against a list of known uncachable files. If 20 the list contains the requested file, the order will be discarded. If the list does not contain the requested file, the order is put on a hold by the Updater 240 so that there will be time for the local cache server 200 to retrieve the file from its origin server.

25 At a time convenient to the central file server 230, i.e. at a time with relative low load on the central server, the central server sends a message to the Updater 240 stating that any pending order should be executed, the reception of this message at the Updater 240 is 30 indicated with step 502. In the next step 503, the execution of the order starts with that the Updater requests a copy of the file, which now should have been retrieved and cached locally, from the local cache server 200 from which the file request originated. A copy of the file is 35 then received from the local cache server in step 504. In step 505, the received file copy is transferred to the central file server 230 to be cached by the same. In the

final step 506, the Updater 240 instructs the Feeder 210 to add the query number corresponding to the file cached at the central file server 230 to the indexed hash table 215.

5 The operation of the central file server 230 is straight forward. Basically it does two things, it answers SQL queries from the Feeders 210 by transferring cached files to them, and, it stores new information files in its cache, which files are transferred to it
10 from the Updater 240.

Another exemplifying embodiment of an Internet caching system according to the present invention is now described with reference to Fig 6. In Fig 6, the system differs from the one shown in Fig 2 in that the system
15 has more than one central file server, here being exemplified with three central cache servers 630. Also, Fig. 6 includes two Feeders 610, each of which is connected to its own set of local cache servers 600. The Feeders 610 and the Updater 640 are arranged together with the
20 central file servers 630 at a central cache site 690. By means of an Ethernet network 680 arranged within the central cache site, the Updater 640 and each Feeder 610 are connected to all central file servers 630.

The additional number of central file servers in
25 this embodiment enables more files to be cached and even more SQL queries to be answered by the central file servers in comparison with the embodiment of Fig. 2. Since the system is completely scaleable, any number of Feeders, Updaters or central file servers can in theory
30 be added to the system.

The basic difference of the operation of the system in Fig. 6 to that of the system in Fig. 2 is that a Feeder 610 need to select one server, out of the plurality of central file servers 630, to which an SQL query
35 should be directed. Each central file server 630 caches information files within original host names within a predefined range. Therefore, the selection of one of the

central file servers is done in accordance with the host name included in the URL received from the local server, either as part of an ICP query or as part of an HTTP request. When one of the central file servers has been 5 selected by the Feeder, the SQL query with the derived query number is directed to that selected file server.

It is understood that the construction and function of the elements described with reference to the drawings will become apparent for those skilled in the art.

10 Even though the invention has been described with reference to specific exemplifying embodiments, many different alterations, modifications and the like will become apparent for those skilled in the art. The described embodiments are therefore not intended to limit 15 the scope of the invention, as defined by the appended claims.

CLAIMS

1. A method for serving requests for Internet information files in an Internet caching system, comprising the steps of:

receiving, at a local Internet cache server, a user request from a user for an Internet information file;

in response to the received request, making a query for said information file, if said information file has not been cached by said local server;

in response to a reply to said query, making a file request for said information file, wherein said file request is directed to a feeder means if said reply indicates that a central file server, storing cached Internet information files, has said information file cached; and

querying, from said feeder means in response to said file request, said central file server for said information file,

in order to decrease the load on said central file server.

2. The method as claimed in claim 1, wherein said query is performed by said local cache server in accordance with a protocol used for communicating between Internet cache servers.

3. The method as claimed in claim 2, wherein said protocol is the Internet Cache Protocol (ICP).

30

4. The method as claimed in claim 2, wherein said protocol is the Cache Digest.

35 5. The method as claimed in any one of claims 1 - 3, wherein said query is directed by said local cache server to said feeder means, which feeder means as a response returns said reply.

6. The method as claimed in 5, comprising the step of deriving, at said feeder means, a query number corresponding to said information file being concerned in said query.

7. The method as claimed in 6, wherein said querying step comprises using the derived query number when querying said central file server for said information file.

8. The method as claimed in claim 6, wherein said query provides an alphanumerical string associated with said information file, said string being used in said step of deriving said query number.

9. The method as claimed in claim 8, wherein said alphanumerical string is a Uniform Resource Locator (URL) and said query number is derived from said URL and at least part of a header information field of said query.

10. The method as claimed in any one of claims 1, 2 or 4, wherein said file request provides an alphanumerical string associated with said information file, said string being used by said feeder means for deriving a query number corresponding to said information file.

11. The method as claimed in claim 10, wherein said alphanumerical string is a Uniform Resource Locator (URL) and said query number is derived from said URL and at least part of a header information field of said file request.

12. The method as claimed in any one of the preceding claims, comprising the step of creating an indexed table having an entry for each Internet information file being cached at said central file server.

13. The method as claimed in claim 12, comprising the steps of:

5 performing a search in said indexed table for said information file; and

indicating in said reply to said query whether or not said information file was found during said search.

14. The method as claimed in any one of the preceding claims, wherein said querying step comprises using the Structured Query Language (SQL) when querying said central file server for said information file.

15. The method as claimed in any one of the preceding claims, wherein said querying step comprises the steps of:

20 selecting, based upon an original host name or IP-address of said information file, a central file server out of a set of central file servers, each server of said set being arranged to cache Internet information files with original host names or IP-addresses within a predefined range; and

querying the selected central file server for said information file.

25

16. The method as claimed in any one of claims 6 - 14, wherein said querying step comprises the steps of:

30 selecting, based upon said query number derived for said information file, a central file server out of a set of central file servers, each server of said set being arranged to cache Internet information files with corresponding query numbers within a predefined range; and

querying the selected central file server for said information file.

35

17. The method as claimed in any one of claims 1 - 16, comprising the further steps of:

retrieving, at said local cache server, said information file from its origin server if said reply to said query indicates that said information file is not cached at said central file server;

5 caching said information file at said local cache server; and

updating said central file server by requesting a copy of said information file from said local cache server and caching said copy in said central file server.

10

18. An arrangement in an Internet caching system, said system comprising at least one local cache server and at least one central file server, both of which servers stores cached Internet information files, which 15 arrangement, for decreasing the load on said central file server, includes a Feeder communicating with said local cache server and with said central file server, wherein said Feeder includes:

first means for receiving a request for an Internet 20 information file from said local cache server;

second means for deriving a query from an alpha-numerical string received from said local cache server; and

third means for querying said central file server 25 for said Internet information file using said query derived by said second means.

19. The arrangement as claimed in claim 18, wherein said first means is arranged to operate in accordance 30 with a layer three Internet protocol.

20. The arrangement as claimed in claim 18 or 19, wherein said third means is arranged to use the Structured Query Language(SQL) when querying for said 35 Internet information file.

21. The arrangement as claimed in any one of claims 18 - 20, wherein said alphanumerical string is included in said request received from said local cache server.

5 22. The arrangement as claimed in claim 21, wherein said query is derived from said alphanumerical string and at least part of a header information field of said request received from said local cache server.

10 23. The arrangement as claimed in claim 22, wherein said query comprises a query number, the query number being derived by applying a hash algorithm to said string and to said part of said header information field.

15 24. The arrangement as claimed in any one of claims 18 - 20, wherein said Feeder includes:
fourth means for receiving a query for an Internet information file from said local cache server; and
fifth means for providing said local cache server
20 with a reply to the received query.

25 25. The arrangement as claimed in claim 24, wherein said fourth means and said fifth means are arranged to operate in accordance with a protocol used for communicating between Internet cache servers.

26. The arrangement as claimed in claim 25, wherein said protocol is the Internet Cache Protocol (ICP).

30 27. The arrangement as claimed in any one of claims 24 - 26, wherein said alphanumerical string is included in said query received from said local cache server.

35 28. The arrangement as claimed in claim 27, wherein said query derived by said second means is derived from said alphanumerical string and at least part of a header

information field of said query received from said local cache server.

29. The arrangement as claimed in claim 28, wherein
5 said query comprises a query number, the query number -
being derived by applying a hash algorithm to said string
and to said part of said header information field.

30. The arrangement as claimed in one of claims 24 -
10 29, wherein said Feeder includes a table with a copy of
the full index of all Internet information files cached
at said central file server.

31. The arrangement as claimed in claim 30, wherein
15 said reply to said received query by said fifth means is
based on the content of said table.

32. The arrangement as claimed in one of claims 18 -
31, wherein said arrangement, for further decreasing the
20 load on said central file server, includes an Updater
communicating with said local cache server and with said
central file server, wherein said Updater includes:

requesting means for requesting a copy of an
Internet information file stored in a local cache server;
25 and

storing means for storing the thereby received copy
in a central file server.

33. The arrangement as claimed in claim 32, wherein
30 said requesting means are arranged to request a copy of
an information file from its origin server, if a local
cache server storing said information file resides behind
a firewall.

35 34. The arrangement as claimed in claim 32 or 33,
wherein said Updater is arranged to communicate with said

Feeder for receiving an order to request said copy of said information file.

35. The arrangement as claimed in any one of claims
5 32 - 34, wherein said Updater includes a list of known -
uncachable information files, for which files a copy
should not be requested.

36. The arrangement as claimed in any one of claims
10 16 - 35, wherein said Feeder is implemented by a lower
end computer and said central file server is implemented
by a higher end computer.

37. The arrangement as claimed in any one of claims
15 32 - 35, wherein said Updater is implemented by a lower
end computer and said central file server is implemented
by a higher end computer.

38. The arrangement as claimed in claim 37, wherein
20 said Updater and at least one Feeder are implemented by a
single lower end computer.

39. An Internet caching system, comprising:
25 a set of local Internet cache servers, wherein each
local cache server is arranged to receive requests from
users for Internet information files;

at least one central file server included in a
central cache site and storing cached Internet informa-
tion files; and

30 feeder means interconnecting said set of local cache
servers with said central file server, said feeder means
including at least one Feeder, which Feeder comprises
means for communicating with at least one local cache
server in accordance with a protocol used for communi-
35 cating between Internet cache servers and means for
retrieving Internet information files from said central

file server using data base queries, thereby decreasing the load on said central file server.

40. The system as claimed in claim 39, wherein said 5 feeder means are included in said central cache site. -

41. The system as claimed in claims 39 or 40, wherein each of said feeder means includes a plurality of Feeders, each of said Feeder interconnecting a subset of 10 said set of local cache servers with said central file server.

42. The Internet caching system as claimed in any one of claims 39 - 41, wherein said central cache site is 15 arranged to serve a defined set of local cache servers, which set in turn serves a linguistically and culturally homogenous user community.

43. The Internet caching system as claimed in any 20 one of claims 39 - 42, wherein said protocol used is either the Internet Cache Protocol or the Cache Digest.



44. The Internet caching system as claimed in any one of claims 39 - 43, wherein each of said Feeder 25 includes a table with a copy of the full index of all information files cached at said central cache site.

45. The Internet caching system as claimed in any one of claims 39 - 44, wherein said central file server 30 includes cached Internet information files having original host names within a predefined range.

46. The Internet caching system as claimed in any one of claims 39 - 45, further comprising updater means, 35 interconnecting said central file server with at least one local cache server of said set, for retrieving a copy of an Internet information file from its origin server or

from said at least one local cache server and for storing said copy in said central file server.

1/5

**FIG. 1****FIG. 2**

2/5




FIG. 3

3/5





FIG. 4

4/5

**FIG. 5**

5/5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US99/21248

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 13/14
 US CL :709/201,218; 711/124,165

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/201,218; 711/124,165

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,787,470 A (DESIMONE et al) 28 JULY 1998 (entire document)	1-46
Y	US 5,774,660 A (BRENDEL et al) 30 JUNE 1998 (entire document)	1-46
A	US 5,794,229 A (FRENCH et al) 11 AUGUST 1998 (entire document)	1-46
A	US 5,809,266 A (TOUMA et al) 15 SEPTEMBER 1998 (entire document)	1-46
A	US 5,633,933 A (AZIZ) 27 MAY 1997 (entire document)	1-46

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	*T*	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X*	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y*	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A*	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search


06 DECEMBER 1999

Date of mailing of the international search report

07 FEB 2000

Name and mailing address of the ISA/US
 Commissioner of Patents and Trademarks
 Box PCT
 Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

 FRANK ASTA

Telephone No. (703) 305-3817